arXiv:1512.08240v1 [stat.ML] 27 Dec 2015

Robust Semi-supervised Least Squares Classification by Implicit Constraints

Jesse H. Krijthe[a,b,*], Marco Loog[a,c]

[a] *Pattern Recognition Laboratory, Delft University of Technology, Mekelweg 4, 2628CD Delft, The Netherlands*
[b] *Department of Molecular Epidemiology, Leiden University Medical Center, Einthovenweg 20, 2333ZC Leiden, The Netherlands*
[c] *Image Group, University of Copenhagen, Universitetsparken 5, DK-2100, Denmark*

Abstract

We introduce the implicitly constrained least squares (ICLS) classifier, a novel semi-supervised version of the least squares classifier. This classifier minimizes the squared loss on the labeled data among the set of parameters implied by all possible labelings of the unlabeled data. Unlike other discriminative semi-supervised methods, this approach does not introduce explicit additional assumptions into the objective function, but leverages implicit assumptions already present in the choice of the supervised least squares classifier. This method can be formulated as a quadratic programming problem and its solution can be found using a simple gradient descent procedure. We prove that, in a limited 1-dimensional setting, this approach never leads to performance worse than the supervised classifier. Experimental results show that also in the general multidimensional case performance improvements can be expected, both in terms of the squared loss that is intrinsic to the classifier, as well as in terms of the expected classification error.

Keywords: Semi-supervised learning, Robust, Least squares classification

1. Introduction

We consider the problem of semi-supervised learning of binary classification functions. As in the supervised paradigm, the goal in semi-supervised learning is to construct a classification rule that maps objects in some input space to a target outcome, such that future objects map to correct target outcomes as well as possible. In the supervised paradigm this mapping is learned using a set of L training objects and their corresponding outputs. In the semi-supervised scenario we are given an additional and often large set of U unlabeled objects.

*Corresponding Author
Email addresses: jkrijthe@gmail.com (Jesse H. Krijthe), m.loog@tudelft.nl (Marco Loog)

The challenge of semi-supervised learning is to incorporate this additional information to improve the classification rule.

The goal of this work is to build a semi-supervised version of the least squares classifier that is robust against deterioration in performance meaning that, at least in expectation, its performance is not worse than supervised least squares classification. While it may seem like an obvious requirement for any semi-supervised method, current approaches to semi-supervised learning do not have this property. In fact, performance can significantly degrade as more unlabeled data is added, as has been shown in [1, 2], among others. This makes it difficult to apply these methods in practice, especially when there is a small amount of labeled data to identify possible reduction in performance. A useful property of any semi-supervised learning procedure would therefore be that its performance does not degrade as we add more unlabeled data. Additionally, many semi-supervised learning procedures are formulated as hard-to-optimize, non-convex objective functions. A more satisfactory state of affairs for semi-supervised classification would therefore be methods that are easier to train and that, on average, do not lead to worse classification performance than their supervised alternatives.

We present a novel approach to semi-supervised learning for the least squares classifier that we will refer to as implicitly constrained least squares classification (ICLS). ICLS leverages implicit assumptions present in the supervised least squares classifier to construct a semi-supervised version. This is done by minimizing the supervised loss function subject to the constraint that the solution has to correspond to the solution of the least squares classifier for some labeling of the unlabeled objects.

As this work is specifically concerned with least squares classification, we note several reasons why this is a particularly interesting classifier to study: First of all, the least squares classifier is a discriminative classifier. Some have claimed semi-supervised learning without additional assumptions is impossible for discriminative classifiers [3, 4]. Our results show this does not strictly hold.

Secondly, the closed-form solution for the supervised least squares classifier allows us to study its theoretical properties. In particular, in the univariate setting without intercept and assuming perfect knowledge of P_X, the distribution of the feature, we show this procedure *never* gives worse performance in terms of the squared loss criterion compared to the supervised least squares classifier. Moreover, using the closed-form solution we can rewrite our semi-supervised approach as a quadratic programming problem, which can be solved through a simple gradient descent with boundary constraints.

Lastly, least squares classification is a useful and adaptable classification technique allowing for straightforward use of, for instance, regularization, sparsity penalties or kernelization [5, 6, 7, 8, 9]. Using these formulations, it has been shown to be competitive with state-of-the-art methods based on loss functions other than the squared loss [7] as well as computationally efficient on large datasets [10].

This work builds on [11] and offers a more complete exposition: we show ICLS can be formulated as a quadratic programming problem, we extend the

experimental results section by including an alternative semi-supervised procedure, adding additional datasets and discussing the peaking phenomenon. Moreover, we extend the theoretical result with conditions when one is likely to see improvement of the proposed approach over the supervised classifier.

The main contributions of this paper are

- A novel convex formulation for robust semi-supervised learning using squared loss (Equation 5)
- A proof that this procedure never reduces performance in terms of the squared loss for the 1-dimensional case without intercept (Theorem 1)
- An empirical evaluation of the properties of this classifier (Section 6)

The rest of this paper is organized as follows. Section 2 gives an overview of related work on semi-supervised learning. Section 3 gives a high level overview of the method while Section 4 introduces our semi-supervised version of the least squares classifier in more detail. We then derive a quadratic programming formulation and present a simple way to solve this problem through bounded gradient descent. Section 5 contains a proof of the improvement of the ICLS classifier over the supervised alternative. This proof is specific to classification with a single feature, without including an intercept in the model. For the multivariate case, we present an empirical evaluation of the proposed approach on benchmark datasets in Section 6 to study its properties. The final sections discuss the results and conclude.

2. Related Work

Many diverse approaches to semi-supervised learning have been proposed [12, 13]. While semi-supervised techniques have shown promise in some applications, such as document classification [14], peptide identification [15] and cancer recurrence prediction [16], it has also been observed that these techniques may give performance worse than their supervised counterparts. See for instance [1, 2], for an analysis of this problem, and [17] for a practical example in part-of-speech tagging. In these cases, disregarding the unlabeled data would lead to better performance.

Some [18, 19] have argued that *agnostic* semi-supervised learning, which [18] defines as semi-supervised learning that is at least no worse than supervised learning, can be achieved by cross-validation on the limited labeled data. Agnostic semi-supervised learning follows if we only use semi-supervised methods when their estimated cross-validation error is significantly lower than those of the supervised alternatives. As the results of [18] indicate, this criterion may be too conservative: given the small amount of labeled data, a semi-supervised method will only be preferred if the difference in performance is very large. If the difference is less distinct, the supervised learner will always be preferred and we potentially ignore useful information from the unlabeled objects. Moreover, this cross-validation approach can be computationally demanding.

Self-Learning

A simple approach to semi-supervised learning is offered by the self-learning procedure [20] also known as Yarowsky's algorithm [21, 22] or retagging [17]. Taking any classifier, we first estimate its parameters on only the labeled data. Using this trained classifier we label the unlabeled objects and add them, or potentially only those we are most confident about, with their predicted labels to the labeled training set. The classifier parameters are re-estimated using these labeled objects to get a new classifier. One iteratively applies this procedure until the predicted labels of the unlabeled data no longer change.

One of the advantages of this procedure is that it can be applied to any supervised classifier. It has also shown practical success in some application domains, particularly document classification [14, 22]. Unfortunately, the process of self-training can also lead to severely decreased performance, compared to the supervised solution [1, 2]. One can imagine that once an object is incorrectly labeled and added to the training set, its incorrect label may be reinforced, leading the solution away from the optimum. Self-learning is closely related to expectation maximization (EM) based approaches [21]. Indeed, expectation maximization suffers from the same issues as self-learning [13]. In Section 6 we compare the proposed approach to self-learning for the least squares classifier.

Additional Assumptions

Some semi-supervised methods leverage the unlabeled data by introducing assumptions that link properties of the features alone to properties of the label of an object given its features. Commonly used assumptions are the smoothness assumption: objects that are close in the feature space likely share the same label; the cluster assumption: objects in the same cluster share a label; and the low density assumption enforcing that the decision boundary should be in a region of low data density.

The low-density assumption is used in entropy regularization [23] as well as for support vector classification in the transductive support vector machine (TSVM) [24] and closely related semi-supervised SVM (S^3VM) [25, 26]. In these approaches an additional term is added to the objective function to push the decision boundary away from regions of high density. Several approaches have been put forth to minimize the resulting non-convex objective function, such as the convex concave procedure [27] and difference convex programming [26, 28].

In all these approaches to semi-supervised learning, a parameter controls the importance of the unlabeled points. When the parameter is correctly set, it is clear, as [19] claims, that TSVM is always no worse than supervised SVM. It is, however, non-trivial to choose this parameter, given that semi-supervised learning is most interesting in cases where we have limited labeled objects, making a choice using cross-validation very unstable. In practice, therefore, TSVM can also lead to performance worse than the supervised support vector machine.

Safe Semi-supervised Learning

[29, 30] attempt to guard against the possibility of deterioration in performance by not introducing additional assumptions, but instead leveraging implicit assumptions already present in the choice of the supervised classifier. These assumptions link parameters estimates that depend on labeled data to parameter estimates that rely on all data. By exploiting these links, semi-supervised versions of the nearest mean classifier and the linear discriminant are derived. Because these links are unique to each classifier, the approach does not generalize directly to other classifiers. The method presented here is similar in spirit, but unlike [29, 30], no explicit equations have to be formulated to link parameter estimates using only labeled data to parameter estimates based on all data. Moreover, our approach allows for theoretical analysis of the non-deterioration of the performance of the procedure.

Aside from the work by [29, 30], another attempt to construct a robust semi-supervised version of a supervised classifier has been made in [31], which introduces the safe semi-supervised support vector machine (S^4VM). This method is an extension of S^3VM [25] which constructs a set of low-density decision boundaries with the help of the additional unlabeled data, and chooses the decision boundary, which, even in the worst-case, gives the highest gain in performance over the supervised solution. If the low-density assumption holds, this procedure provably increases classification accuracy over the supervised solution. The main difference with the method considered in this paper, however, is that we make no such additional assumptions. We show that even without these assumptions, safe improvements are possible for the least squares classifier.

Semi-supervised Least Squares

While least squares classification has been widely used and studied [5, 6, 8], little work has been done on applying semi-supervised learning to the least squares classifier specifically. For least squares regression, [32] describe an iterative method for handling missing outcomes that was formally proposed in [33]. In the case of least squares regression, this method has some computational advantages over discarding the unlabeled data but its solution always coincides with the supervised solution. [34] studied the value of knowing $\mathbb{E}[\mathbf{X}^\top\mathbf{X}]$, where $\mathbf{X}$ is the $L \times d$ design matrix containing the feature values for each observation. If we assume the number of unlabeled data points is large, this is similar to the semi-supervised situation. It is shown that if the size of the parameters is small compared to the noise, the variance of a procedure that plugs in $\mathbb{E}[\mathbf{X}^\top\mathbf{X}]$ as the estimate of $\mathbf{X}^\top\mathbf{X}$ has a lower variance than supervised least squares regression. As the size of the parameters increases, this effect reverses. In fact, the paper demonstrates that in this semi-supervised setting no best linear unbiased estimator for the regression coefficients exists. In Section 6, we compare our approach to using this plug-in estimate by substituting the matrix $\mathbf{X}^\top\mathbf{X}$ by a version based on both labeled and unlabeled data. A similar plug-in procedure has been used by [35] for linear discriminant analysis for dimensionality reduction which is closely related to least squares classification. Here the (normalized) total scatter matrix, which plays a similar role to the $\mathbf{X}^\top\mathbf{X}$ matrix

in least squares regression is exchanged with the more accurate estimate of the total scatter based on both labeled and unlabeled data.

3. Implicitly Constrained Least Squares Classification

Given a limited set of L labeled objects and a potentially large set of U unlabeled objects, the goal of implicitly constrained least squares classification is to use the latter to improve the solution of the least squares classifier trained on just the labeled data. We start with a sketch of this approach, before discussing the details.

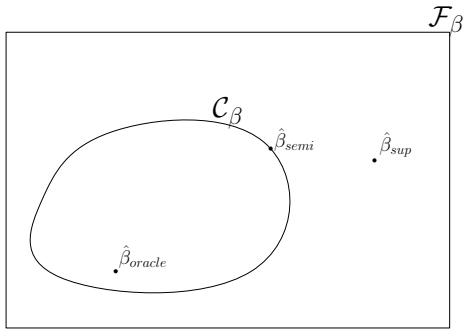

Figure 1: A visual representation of implicitly constrained semi-supervised learning. $\mathcal{F}_\beta$ is the space of all linear models. $\hat{\beta}_{sup}$ denotes the solution given only a small amount of labeled data. $\mathcal{C}_\beta$ is the subset of the space which contains all the solutions we get when applying all possible (soft) labelings to the unlabeled data. $\hat{\beta}_{semi}$ is a projection of $\hat{\beta}_{sup}$ onto $\mathcal{C}_\beta$. $\hat{\beta}_{oracle}$ is the supervised solution if we would have the labels for all the objects.

Given the supervised least squares classifier, consider the hypothesis space of all possible parameter vectors, which we will denote as $\mathcal{F}_\beta$, see Figure 1. Given a set of labeled objects, we can determine the supervised parameter vector $\hat{\beta}_{sup}$. Suppose we also have a potentially large number U of unlabeled objects. Assume that every object has a label, it is merely unknown to us. If these labels were to be revealed, it is clear how the additional objects can improve classification performance: we estimate the least squares classifier using all the data to obtain the parameter vector $\hat{\beta}_{oracle}$. Since this estimate is based on more objects, we expect the parameter estimate to be better. These real labels are unknown, but we can still consider all possible labelings of unlabeled objects, and estimate corresponding parameters based on these imputed labelings. In this way, we get a set of possible parameters for our classifier, which form the set denoted by $\mathcal{C}_\beta \subset \mathcal{F}_\beta$. Clearly one of these labelings corresponds to the real, but unknown, labeling, so one of the parameter estimates in this set corresponds to the solution we would obtain using all the correct labels of both the labeled and unlabeled objects. Because these are the only possible classifiers when the true labels would be revealed, we propose to look within this set $\mathcal{C}_\beta$ for an improved semi-supervised solution.

Two issues then remain: how do we choose the best parameters from this set and how do we find these without having to enumerate all possible labelings?

Looking at the first problem, we reiterate that the goal of semi-supervised learning is to find a good classification rule and, therefore, still the obvious way to evaluate this rule is by the loss on the labeled training points. In other words, we choose the classifier from the parameter set that minimizes the squared loss on the labeled points. Note this approach is rather different from other approaches to semi-supervised learning where the loss is adapted by including a term that depends on the unlabeled data points. In our formulation, the loss function is still the regular, supervised loss of our classification procedure. We will denote this solution by $\hat{\beta}_{semi}$.

As for the second issue, after relaxing the constraint that we need hard labels for the data points, we will see that the resulting optimization problem is, in fact, an instantiation of well-studied quadratic programming, which we solve using a simple gradient descent procedure.

4. Method

4.1. Supervised Multivariate Least Squares Classification

Least squares classification [5, 7] is the direct application of well-known ordinary least squares regression to a classification problem. A linear model is assumed and the parameters are minimized under squared loss. Let $\mathbf{X}$ be an $N_l \times (m+1)$ design matrix with N_l rows containing vectors of length equal to the number of features m plus a constant feature to encode the intercept. Vector $\mathbf{y}$ denotes an $N_l \times 1$ vector of class labels. We encode one class as 0 and the other as 1. The multivariate version of the empirical risk function for least squares estimation is given by

$$\hat{R}(\boldsymbol{\beta}) = \frac{1}{n} \left\| \mathbf{X}\boldsymbol{\beta} - \mathbf{y} \right\|_2^2 . \tag{1}$$

The well-known closed-form solution for this problem is found by setting the derivative with respect to $\boldsymbol{\beta}$ equal to $\mathbf{0}$ and solving for $\boldsymbol{\beta}$, giving

$$\hat{\boldsymbol{\beta}} = \left(\mathbf{X}^T\mathbf{X}\right)^{-1} \mathbf{X}^T\mathbf{y} . \tag{2}$$

In case $\mathbf{X}^T\mathbf{X}$ is not invertible (for instance when $n < (d+1)$), a pseudo-inverse is applied. As we will see, the closed form solution to this problem will enable us to formulate our semi-supervised learning approach in terms of a standard quadratic programming problem, which is easy to optimize.

4.2. Implicitly Constrained Least Squares Classification

In the semi-supervised setting, apart from a design matrix $\mathbf{X}$ and target vector $\mathbf{y}$, an additional set of measurements $\mathbf{X}_u$ of size $U \times (d+1)$ *without* a corresponding target vector $\mathbf{y}_u$ is given. In what follows, $\mathbf{X}_e = \begin{bmatrix} \mathbf{X}^T & \mathbf{X}_u^T \end{bmatrix}^T$

denotes the extended design matrix which is simply the concatenation of the design matrices of the labeled and unlabeled objects.

In the implicitly constrained approach, we incorporate the additional information from the unlabeled objects by searching within the set of classifiers that can be obtained by all possible labelings $\mathbf{y}_u$, for the one classifier that minimizes the *supervised* empirical risk function (1). This set, $\mathcal{C}_\beta$, is formed by the $\boldsymbol{\beta}$s that would follow from training supervised classifiers on all (labeled and unlabeled) objects going through all possible soft labelings for the unlabeled samples, i.e., using all $\mathbf{y}_u \in [0,1]^U$. Since these supervised solutions have a closed form, this can be written as

$$\mathcal{C}_\beta := \left\{ \boldsymbol{\beta} = \left(\mathbf{X}_e^\top \mathbf{X}_e\right)^{-1} \mathbf{X}_e^\top \begin{bmatrix} \mathbf{y} \\ \mathbf{y}_u \end{bmatrix} : \mathbf{y}_u \in [0,1]^U \right\} . \tag{3}$$

The soft labeling provides both a relaxation for computational reasons as well as a strategy to deal with label uncertainty. We can interpret these fractions as a type of class posterior for the unlabeled objects. This constraint set $\mathcal{C}_\beta$, combined with the supervised loss that we want to optimize in Equation (1), gives the following definition for implicitly constrained semi-supervised least squares classification:

$$\underset{\boldsymbol{\beta} \in \mathcal{C}_\beta}{\operatorname{argmin}} \quad \hat{R}(\boldsymbol{\beta}) . \tag{4}$$

Since $\boldsymbol{\beta}$ is fixed for a particular choice of $\mathbf{y}_u$ and has a closed form solution, we can rewrite the minimization problem in terms of $\mathbf{y}_u$ instead of $\boldsymbol{\beta}$:

$$\underset{\mathbf{y}_u \in [0,1]^U}{\operatorname{argmin}} \quad \frac{1}{n} \left\| \mathbf{X} \left(\mathbf{X}_e^\top \mathbf{X}_e\right)^{-1} \mathbf{X}_e^\top \begin{bmatrix} \mathbf{y} \\ \mathbf{y}_u \end{bmatrix} - \mathbf{y} \right\|_2^2 . \tag{5}$$

The problem defined in Equation (5) can be written in a standard quadratic programming form:

$$\begin{aligned} &\min_{\mathbf{y}_u} \frac{1}{2} \mathbf{y}_u^T \mathbf{Q} \mathbf{y}_u + \mathbf{c}^T \mathbf{y}_u \\ &\text{subject to:} \quad \begin{bmatrix} \mathbf{I}_U \\ -\mathbf{I}_U \end{bmatrix} \mathbf{y}_u \leq \begin{bmatrix} \mathbf{1}_U \\ \mathbf{0}_U \end{bmatrix} \end{aligned} \tag{6}$$

where

$$\mathbf{Q} = \frac{1}{n} \mathbf{X}_u \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{X} \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}_u^T$$

and

$$\begin{aligned} \mathbf{c} = &\frac{1}{n} \mathbf{X}_u \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{y} \\ &+ \frac{1}{2n} \mathbf{X}_u \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{X} \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{y} . \end{aligned}$$

Where $\mathbf{I}_U$ denotes the $U \times U$ identity matrix and $\mathbf{1}_U$ and $\mathbf{0}_U$ denote column vectors of respectively ones and zeros.

Since the matrix $\mathbf{Q}$ is a product of a matrix and its transpose, it is guaranteed to be positive semi-definite. The problem is typically not positive definite

because there are different labelings that will lead to one and the same minimum objective.

The quadratic problem defined above can be solved using, for instance, an interior point method. We have found a gradient descent approach to be easier to apply. Taking the derivative with respect to $\mathbf{y}_u$ and rearranging the terms we find

$$\begin{aligned}\frac{\partial \hat{R}}{\partial \mathbf{y}_u} = &\frac{2}{n}\mathbf{X}_u \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{X} \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{y} \\ &+ \frac{2}{n}\mathbf{X}_u \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{X} \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}_u^T \mathbf{y}_u \\ &+ \frac{2}{n}\mathbf{X}_u \left(\mathbf{X}_e^T \mathbf{X}_e\right)^{-1} \mathbf{X}^T \mathbf{y} \,.\end{aligned}$$

Because of its convexity, this problem can be solved efficiently using a quasi-Newton approach that allows for the $[0, 1]$ box bounds, such as L-BFGS-B [36]. Solving for $\mathbf{y}_u$ gives a labeling that we can use to construct the semi-supervised classifier using Equation (2) by considering the imputed labels as the labels for the unlabeled data.

5. Theoretical Results

We will examine this procedure by considering it in a limited, yet illustrative setting. In this case we will, in fact, prove that our procedure will *never* give a worse least squares estimate than the supervised solution Consider the case where we have just one feature x, a limited set of labeled instances and assume we know the probability density function of this feature f_X exactly. This last assumption is similar to having unlimited unlabeled data and is also considered, for instance, in [37]. We consider a linear model with no intercept: $y = x\beta$ where y, without loss of generality, is set as 0 for one class and 1 for the other. For new data points, estimates $\hat{y}$ can be used to determine the predicted label of an object by using a threshold set at, for instance, 0.5.

The expected squared loss, or risk, for this model is given by

$$R^*(\beta) = \sum_{y \in \{0,1\}} \int_{-\infty}^{\infty} (x\beta - y)^2 f_{X,Y}(x, y) \mathrm{d}x \,. \tag{7}$$

Where $f_{X,Y} = P(y|x) f_X(x)$. We will refer to this as the joint density of X and Y. Note, however, that this is not strictly a density, since it deals with the joint distribution over a continuous X and a discrete Y. The optimal solution β^* is given by the β that minimizes this risk

$$\beta^* = \operatorname*{argmin}_{\beta \in \mathbb{R}} R^*(\beta) \,. \tag{8}$$

We will show the following result:

Theorem 1. *Given a linear model in 1D without intercept, $y = x\beta$, and f_X known, the estimate obtained through implicitly constrained least squares always has an equal or lower risk than the supervised solution*

$$R^*(\hat{\beta}_{semi}) \leq R^*(\hat{\beta}_{sup})\,.$$

In particular, given 1 labeled sample, if $f_{X,Y}$ is continuous in the feature X with bounded second moment and $f_{X,Y}(0,1) > 0$, then

$$\mathbb{E}[R^*(\hat{\beta}_{semi})] < \mathbb{E}[R^*(\hat{\beta}_{sup})]\,.$$

Proof. Setting the derivative of (7) with respect to β to 0 and rearranging we get

$$\begin{aligned}
\beta &= \left(\int_{-\infty}^{\infty} x^2 f_X(x)\mathrm{d}x\right)^{-1} \sum_{y\in\{0,1\}} \int_{-\infty}^{\infty} xy f_{X,Y}(x,y)\mathrm{d}x && (9)\\
&= \left(\int_{-\infty}^{\infty} x^2 f_X(x)\mathrm{d}x\right)^{-1} \int_{-\infty}^{\infty} x f_X(x) \sum_{y\in\{0,1\}} yP(y|x)\mathrm{d}x && (10)\\
&= \left(\int_{-\infty}^{\infty} x^2 f_X(x)\mathrm{d}x\right)^{-1} \int_{-\infty}^{\infty} x f_X(x)\mathbb{E}[y|x]\mathrm{d}x\,. && (11)
\end{aligned}$$

In this last equation, since we assume $f_X(x)$ as given, the only unknown is the function $\mathbb{E}[y|x]$, the expectation of the label y, given x. Now suppose we consider every possible labeling of the unlimited number of unlabeled objects including fractional labels, that is, every possible function where $\mathbb{E}[y|x] \in [0,1]$. Given this restriction on $\mathbb{E}[y|x]$, the second integral in (11) becomes a re-weighted version of the expectation operation over x. By changing the choice of $\mathbb{E}[y|x]$ one can vary the value of this integral, but it will always be bounded on an interval on $\mathbb{R}$. It follows that all possible β's also form an interval on $\mathbb{R}$, which is the constraint set $\mathcal{C}_\beta$. The optimal solution has to be in this interval, since it corresponds to a particular but unknown $\mathbb{E}[y|x]$.

Using the set of labeled data, we can construct a supervised solution $\hat{\beta}_{sup}$ that minimizes the loss on the training set of L labeled objects (see Figure 2)

$$\hat{\beta}_{sup} = \operatorname*{argmin}_{\beta\in\mathbb{R}} \sum_{i=1}^{L} (x_i\beta - y_i)^2\,. \tag{12}$$

Now, either this solution falls within the constrained region, $\hat{\beta}_{sup} \in \mathcal{C}_\beta$ or not, $\hat{\beta}_{sup} \notin \mathcal{C}_\beta$, with different consequences:

1. If $\hat{\beta}_{sup} \in \mathcal{C}_\beta$ there is a labeling of the unlabeled points that gives us the same value for β. Therefore, the solution falls within the allowed region and there is no reason to update our estimate. Therefore $\hat{\beta}_{semi} = \hat{\beta}_{sup}$.

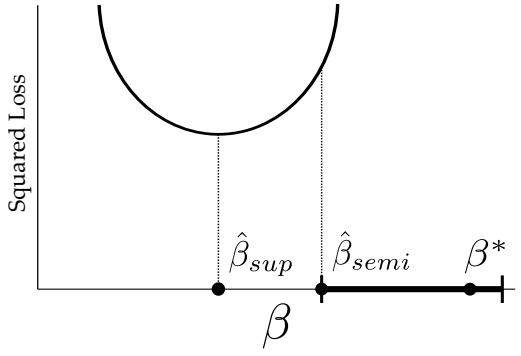

Figure 2: An example where implicitly constrained optimization improves performance. The supervised solution $\hat{\beta}_{sup}$ which minimizes the supervised loss (the solid curve), is not part of the interval of allowed solutions. The solution that minimizes this supervised loss within the allowed interval is $\hat{\beta}_{semi}$. This solution is closer to the optimal solution β^* than the supervised solution $\hat{\beta}_{sup}$.

2. Alternatively, if $\hat{\beta}_{sup} \notin \mathcal{C}_\beta$, the solution is outside of the constrained region (as shown in Figure 2): there is no possible labeling of the unlabeled data that will give the same solution as $\hat{\beta}_{sup}$. We then update the β to be the β within the constrained region that minimizes the loss on the supervised training set. As can be seen from Figure 2, this will be a point on the boundary of the interval. Note that $\hat{\beta}_{semi}$ is now closer to β^* than $\hat{\beta}_{sup}$. Since the true loss function $R^*(\beta)$ is convex and achieves its minimum in the optimal solution, corresponding to the true labeling, the risk of our semi-supervised solution will always be equal to or lower than the loss of the supervised solution.

Thus, the proposed update either improves the estimate of the parameter β or it does not change the supervised estimate. In no case will the semi-supervised solution be worse than the supervised solution, in terms of the expected squared loss. This concludes the proof of the first part of the theorem.

The last part of the theorem gives a general condition when, in expectation, our semi-supervised approach will outperform the supervised learner. Because $\hat{\beta}_{semi}$ will never be worse than $\hat{\beta}_{sup}$, to prove this we only need to show that for some observation of a labeled point with positive $f_{X,Y}(x, y) > 0$, the estimated $\hat{\beta}_{sup}$ is outside of the interval $\mathcal{C}_\beta$, in which case $R^*(\hat{\beta}_{semi}) < R^*(\hat{\beta}_{sup})$.

If we observe an object labeled 1 with feature value x, the corresponding estimate $\hat{\beta}_{sup} = \frac{1}{x}$. Since the improvement in loss will only result if this estimate is not in the constrained region, we need to show that

$$P(\tfrac{1}{x} \notin \mathcal{C}_\beta, y = 1) > 0\,. \tag{13}$$

To do this, consider the bounds of the interval $\mathcal{C}_\beta$. These most extreme values are obtained whenever all negative values of x are assigned label 0 while

the positive x get labels 1, or the other way around. From (11) and writing $\mathbb{E}(X^2) = \int_{-\infty}^{\infty} x^2 f_X(x)\mathrm{d}x$ we find the interval is given by

$$\mathcal{C}_\beta = \left[\frac{\int_{-\infty}^{0} x f_X(x)\mathrm{d}x}{\mathbb{E}(X^2)}, \frac{\int_{0}^{\infty} x f_X(x)\mathrm{d}x}{\mathbb{E}(X^2)} \right] . \tag{14}$$

Combining this with (13), we get the condition

$$P\left(\frac{\mathbb{E}(X^2)}{\int_{-\infty}^{0} x f_X(x)\mathrm{d}x} < x < 0 \vee 0 < x < \frac{\mathbb{E}(X^2)}{\int_{0}^{\infty} x f_X(x)\mathrm{d}x}, y = 1 \right) > 0 . \tag{15}$$

Since $f_{X,Y}$ is assumed to be continuous, $\mathbb{E}[X^2] > 0$, and the lower bound in this equation is always smaller than 0, while the upper bound is always larger than 0. The assumption of the continuity of $f_{X,Y}$ ensures that (15) holds whenever $f_{X,Y}(0,1) > 0$. The property $f_{X,Y}(0,1) > 0$ is satisfied by many distributions of the data. The result, therefore, indicates, that in the case of 1 labeled sample improvement is not only possible, but will occur in many cases. When we have multiple labeled examples, this effect will likely become smaller. This makes sense: the more labeled data we have to estimate the parameter, the smaller the impact of the unlabeled objects will be. □

6. Empirical Results

To study the properties of the proposed semi-supervised approach to least squares classification, we compare how, for different sizes of the unlabeled sample, this approach fares against supervised least squares classification without the constraints.

For comparison we include two alternative semi-supervised approaches and an oracle solution:

Self-Learning. Using a simple procedure proposed by [20], among others, the supervised least squares classifier is updated iteratively by using its class predictions on the unlabeled objects as the labels for the unlabeled objects in the next iteration. This is done until convergence.

Updated Second Moment Least Squares (USM). In this approach we replace the second moment matrix $\mathbf{X}^T\mathbf{X}$ with an appropriately scaled matrix $\mathbf{X}_e^\top \mathbf{X}_e$ similar to the estimator studied in [34]:

$$\hat{\boldsymbol{\beta}}_{USM} = \left(\tfrac{L}{L+U} \mathbf{X}_e^\top \mathbf{X}_e \right)^{-1} \mathbf{X}^T \mathbf{y}$$

where $\mathbf{X}_e$ and $\mathbf{y}$ are centered. This ensures that results do not depend on the particular encoding of the labels used. We will refer to this as updated second moment least squares (USM) classification.

Oracle. The performance of the least squares classifier if all unlabeled objects were labeled as well. This serves as the unattainable upper bound on the performance of any semi-supervised learner.

A description of the datasets used for our experiments is given in Table 1. We use datasets from both the UCI repository [38] and from the benchmark datasets proposed by [12]. While the benchmark datasets proposed in [12] are useful, in our experience, the results on these datasets are very homogeneous because of the similarity in their dimensionality and their low Bayes errors. The UCI datasets are more diverse both in terms of the number of objects and features as well as the nature of the underlying problems. Taken together, this collection allows us to investigate the properties of our approach for a wide range of problems. All the code used to run the experiments is available from the first author's website.

Table 1: Description of the datasets used in the experiments. PCA99 refers to the number of principal components required to retain at least 99% of the variance. Majority refers to the proportion of the number of objects from the largest class

Dataset	Objects	Features	PCA99	Majority	Source
Haberman	306	3	3	0.74	[38]
Ionosphere	351	33	30	0.64	[38]
Parkinsons	195	22	12	0.75	[38]
Diabetes	768	8	8	0.65	[38]
Sonar	208	60	43	0.53	[38]
SPECT	267	22	21	0.79	[38]
SPECTF	267	44	37	0.79	[38]
Transfusion	748	4	3	0.76	[38]
WDBC	569	30	17	0.63	[38]
Mammography	961	9	9	0.54	[38]
Digit1	1500	241	221	0.51	[12]
USPS	1500	241	183	0.80	[12]
COIL2	1500	241	114	0.50	[12]
BCI	400	117	45	0.50	[12]
g241c	1500	241	235	0.50	[12]
g241d	1500	241	235	0.50	[12]

6.1. Peaking Behaviour in Semi-supervised Least Squares

With fewer than d samples, the *supervised* least squares classifier that utilizes a pseudo-inverse is known to exhibit a peaking phenomenon, as described in [39, 40]: Starting from a single observation, expected classification errors generally decrease as we add more data before errors increase again to reach a maximum approximately when the number of features is equal to the number of observations. This phenomenon can also be observed in the semi-supervised setting. Figure 3 shows learning curves of the methods considered here, using 10 labeled training objects and and increasing number of unlabeled objects.

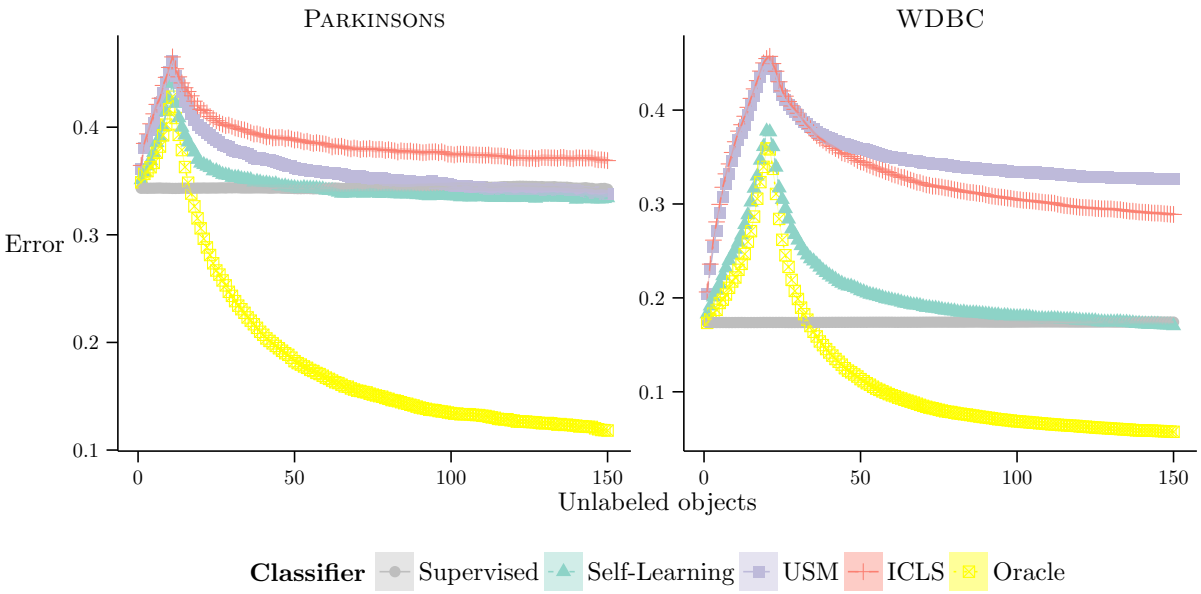

Figure 3: Peaking phenomenon in Semi-supervised Least Squares Classification. The lines indicate the Mean classification error for $L = \max(d + 5, 20)$ and 1000 repeats. The shaded areas indicate $+/-$ the standard error of the mean.

Performance is evaluated on objects that were not in the labeled or unlabeled set. The Oracle classifier indicates the mean error when we do have the labels for the unlabeled objects and therefore corresponds to the peaking phenomenon in the supervised case. In the supervised case, several proposals have been done to ameliorate this peaking behaviour, such as feature selection, regularization, removing objects, injecting noise in the features, or adding redundant features [41]. The semi-supervised learners suffer from the same peaking phenomenon, except that unlike the Oracle, USM and ICLS do not fully recover from the initial increase in classification error.

We have no full explanation for the observed peaking behaviour in the semi-supervised setting. Even in the supervised setting the behaviour remains elusive. The two observation we do make are that the peak occurs at the same location for both the supervised and semi-supervised scenarios and that the subspace defined by the input data is the defining characteristic for the location of the peak.

This peaking behaviour is not the primary topic of this work and in the remainder we will restrict our attention to the case where there are enough labeled objects such that the matrix $\mathbf{X}^\top \mathbf{X}$ is invertible.

6.2. Comparison of Learning Curves

We study the behavior of the expected classification error of the ICLS procedure for different sizes of the unlabeled set. This statistic has two desired properties. First of all it should never be higher than the expected classification error of the supervised solution, which is based on only the labeled data. Secondly, the expected classification error should not increase as we add more unlabeled data. A semi-supervised classifier that has both these properties can be used safely, since adding unlabeled data and continuing to add more unlabeled data will never decrease performance, on average.

Experiments were conducted as follows. For each dataset, L labeled points were randomly chosen, where we make sure to sample at least 1 object from each of the two classes. Since the peaking phenomenon described in the previous section is not main topic of this work, we avoid this situation by considering the setting in which the labeled design matrix is of full rank, which we ensure by setting $L = d + 5$, the dimensionality and intercept of the dataset plus five observations. For all datasets we ensure a minimum of $L = 20$ labeled objects.

Next, we create unlabeled subsets of increasing size $U = [2, 4, 8, ..., 1024]$ by randomly selecting points from the original dataset without replacement. The classifiers are trained using these subsets and the classification performance is evaluated on the remaining objects. Since the test set decreases in size as the number of unlabeled objects increases, the standard error slightly increases with the number of unlabeled objects.

The results of these experiments are shown in Figure 4. We report the mean classification error as well as the standard error of this mean. As can be seen from the tight confidence bands, this offers an accurate estimate of the expected classification error.

This procedure of sampling labeled and unlabeled points is repeated 1000 times and the average classification error (Figure 4) and squared loss (Figure 5) on the test set is determined. The latter is done to evaluate whether the approach is effective in increasing generalization performance in terms of the loss used in estimating the classifier. This is the same loss that we consider in Theorem 1. Even though in applications the ultimate goal may typically be classification performance, this allows us to study whether problems occur because of the optimization itself, or because of the link between the surrogate loss used and the classification error.

We find that, generally, the ICLS procedure has monotonically decreasing error curves as the number of unlabeled samples increases, unlike self-learning. On the Diabetes and Transfusion datasets, the performance of self-learning becomes worse than the supervised solution when more unlabeled data is added, while the ICLS classifier again exhibits a monotonic decrease of the average error rate. The USM classifier performs well on most datasets except for the Mammography dataset, where both in terms of average error rates and squared loss, performance is worse than the supervised classifier.

When we compare the error curves and the loss curves, the non-monotonically decreasing losses for the self-learner correspond to increased errors. In general,

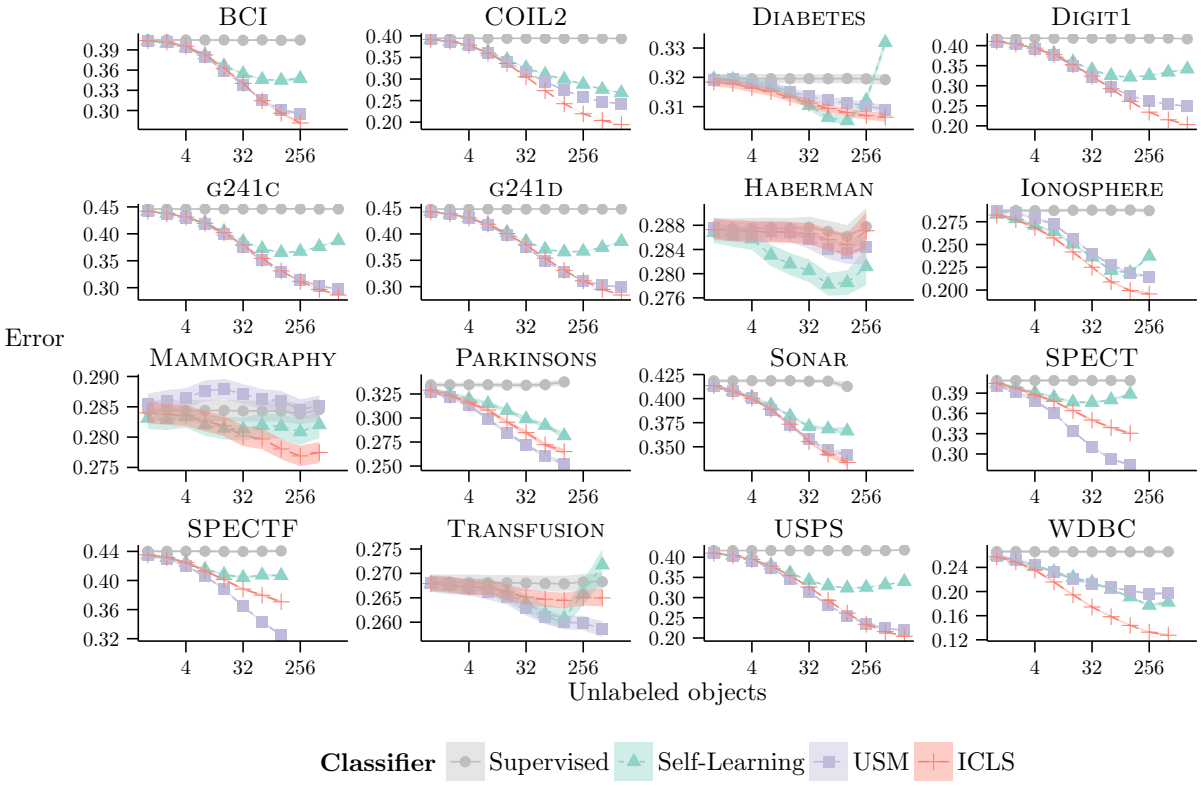

Figure 4: Mean classification error for $L = \max(d+5, 20)$ and 1000 repeats. The shaded areas indicate +/− the standard error of the mean.

however, similar losses for different classifiers can give rise to different behaviours in terms of error rates.

6.3. Benchmark performance

We now consider the performance of these classifiers in a cross-validation setting. The experiment is set up as follows. For each dataset, the objects are randomly divided into 10 folds. We iteratively go through the folds using 1 fold as validation set, and the other 9 as the training set. From this training set, we then randomly select $L = \max(d + 5, 20)$ labeled objects, as in the previous experiment, and use the rest as unlabeled data. After predicting labels for the validation set for each fold, the classification error is then determined by comparing the predicted labels to the real labels. This is repeated 100 times, while randomly assigning objects to folds in each iteration.

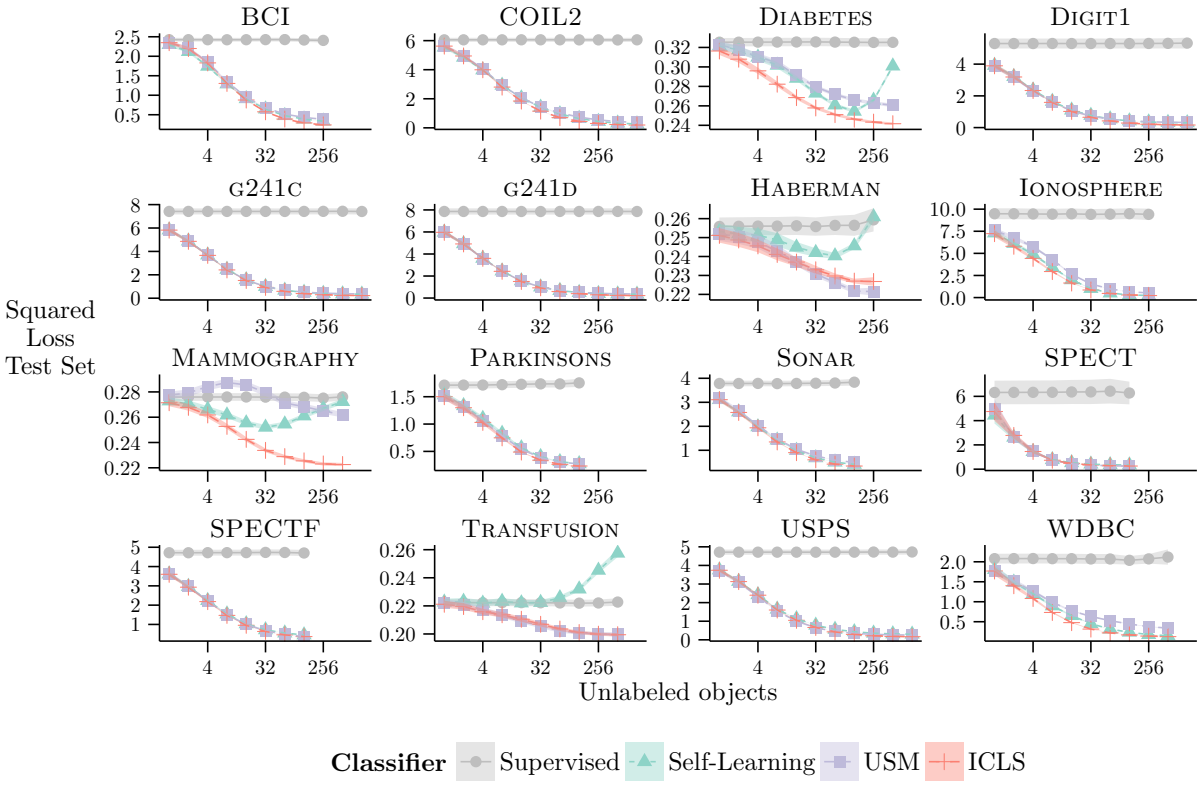

Figure 5: Mean squared loss on the test set for $L = \max(d + 5, 20)$ and 1000 repeats. The shaded areas indicate $+/-$ the standard error of the mean.

The cross-validation procedure used here is slightly different from that described in [12], to make it more closely relate to the cross-validation procedure that is usually employed in supervised learning. More specifically, our procedure ensures the validation sets are independent (non-overlapping), such that, after going over all the folds, each object is in the validation set only once. This is different from the procedure in [12], were the authors ensure the *labeled* sets are non-overlapping. We have not found a qualitative difference in the error rates, however, when using the procedure proposed in [12]. The advantage of the procedure employed here is that every object gets a single predicted label, allowing for the direct comparison of predictions of different classifiers.

The results shown in Table 2 tell a similar story to those in the previous experiment. Most importantly for the purposes of this paper, ICLS, in general, offers solutions that give at least no higher expected classification error than the

Table 2: Average 10-fold cross-validation error and (between parentheses) number of times the error of the semi-supervised classifier is higher than the supervised error for 100 repeats. Indicated in **bold** is which semi-supervised classifier has lowest average error. A Wilcoxon signed rank test at 0.01 significance level is done to determine whether a semi-supervised classifier is significantly worse than the supervised classifier, indicated by underlined values.

Dataset	Supervised	Self-Learning	USM	ICLS	Oracle
HABERMAN	0.29	**0.28 (33)**	0.28 (42)	0.29 (24)	0.26 (11)
IONOSPHERE	0.29	0.24 (1)	0.22 (1)	**0.19 (0)**	0.13 (0)
PARKINSONS	0.34	0.29 (5)	**0.25 (3)**	0.26 (1)	0.12 (0)
DIABETES	0.32	0.34 (83)	0.31 (31)	**0.31 (7)**	0.23 (0)
SONAR	0.42	0.37 (5)	0.34 (3)	**0.33 (1)**	0.25 (0)
SPECT	0.41	0.39 (28)	**0.28 (0)**	0.33 (1)	0.18 (0)
SPECTF	0.43	0.40 (14)	**0.31 (0)**	0.36 (2)	0.23 (0)
TRANSFUSION	0.27	0.28 (63)	**0.26 (30)**	0.27 (25)	0.23 (2)
WDBC	0.27	0.18 (0)	0.20 (2)	**0.13 (0)**	0.04 (0)
MAMMOGRAPHY	0.28	0.28 (28)	0.28 (54)	**0.27 (14)**	0.20 (0)
DIGIT1	0.42	0.34 (0)	0.25 (0)	**0.20 (0)**	0.06 (0)
USPS	0.42	0.34 (0)	0.22 (0)	**0.20 (0)**	0.09 (0)
COIL2	0.39	0.27 (0)	0.24 (0)	**0.19 (0)**	0.10 (0)
BCI	0.41	0.35 (1)	0.30 (0)	**0.28 (0)**	0.16 (0)
G241C	0.45	0.39 (0)	0.30 (0)	**0.29 (0)**	0.14 (0)
G241D	0.45	0.39 (0)	0.30 (0)	**0.29 (0)**	0.13 (0)

supervised procedure. On many of these datasets, the self-learning approach seems to share this property. However, if we look at for how many of the cross-validation repeats the ICLS and self-learning give lower error than the supervised solution, there is a clear difference. The self-learning solution gives a higher error on more of the repeats than ICLS, for all of the datasets.

The results also show that unlabeled information is of use. Particularly on the last six datasets, ICLS and USM offers large improvement in classification accuracy over the supervised solution. The differences in performance between ICLS and self-learning can also be quite substantial, where ICLS outperforms self-learning on most of the datasets. USM performs well on many of the datasets, especially when we consider how simple and computationally efficient this procedure is.

7. Discussion

From Theory to Empirical Results

The results presented in this paper are rather promising, especially in the light of the negative theoretical performance results presented in the literature [1]. The result in Theorem 1, to start with, indicates the proposed procedure is in some way robust against reduction in performance. The strong result of this theorem, stating that performance never gets worse, holds in the 1D case with unlimited unlabeled data and no intercept in the model. A slightly weaker

result, that performance does not degrade on average may still hold without these assumptions. This last statement is corroborated by the empirical results showing improvements in averaged squared errors for ICLS throughout.

The results in the previous section also indicate that such improved results hold in terms of the misclassification error, at least on this collection of datasets. These empirical observations are encouraging because we are often interested in misclassification error and not the squared loss that was considered in Theorem 1. Furthermore the experiments were carried out in the multivariate setting with an intercept term using limited unlabeled data, rather than the unlimited unlabeled data setting considered in the theorem. This indicates that minimizing the supervised loss over the subset $\mathcal{C}_\beta$, leads to a semi-supervised learner with desirable behavior, both theoretically in terms of risk and empirically in terms of classification error.

Robustness

The method considered in this work is different from most previous work in semi-supervised learning in that it is inherently robust against a decrease in performance. The robustness of the method comes from the fact that we do not accept solutions that do not work on the labeled data. The goal of semi-supervised learning is to improve supervised techniques using the additional information inherent in the additional unlabeled objects. Previous approaches have done this by changing the loss function that is being optimized, in particular by introducing an extra term corresponding to assumptions about the unlabeled data. The loss function then becomes a mixture between the supervised objective and an unsupervised objective. If the goal is classification, however, we propose that the loss function should remain the supervised loss function. The unlabeled objects are merely used to introduce constraints on the possible solutions to this loss function, but do not change its functional form.

Assumptions

Most other semi-supervised techniques rely on introducing useful assumptions that link information about the distribution of the features P_X to the posterior of the classes $P_{Y|X}$. It has been argued that, for discriminative classifiers, semi-supervised learning is impossible without these additional assumptions about the link between labeled and unlabeled objects [3, 4]. ICLS, however, is both a discriminative classifier and no explicit additional assumptions about this link are made. Any assumptions that are present follow, implicitly, from the choice of squared loss as the loss function and from the chosen hypothesis space.

In fact, additional assumptions may actually be at the root of the problem: clearly if such an additional assumption is correct, a semi-supervised classifier can gain from it, but if the assumption is incorrect, degraded performance may ensue. What we leverage in our approach are the implicit assumptions that are, in a sense, intrinsic to the supervised least squares classifier.

One could argue that constraining the solutions to $\mathcal{C}_\beta$ is an assumption as well. It corresponds to a very weak assumption about the supervised classifier:

that it will improve when we add additional labeled data. This is generally assumed in the supervised setting as well. The lack of additional assumptions has another advantage: no additional hyperparameter value needs to be selected that the importance of the unlabeled data for the results in Sections 5 and 6 to hold as ICLS acts as a type of data dependent regularization.

Note that the solution provided by self-learning is, by construction, also in the constrained subset $\mathcal{C}_\beta$. The difference with ICLS is that in ICLS the choice of estimate from $\mathcal{C}_\beta$ is based on information of the labeled objects only, while self-learning also uses the imputed labels on the unlabeled objects. This may lead to self-deception: if the imputed labels are wrong, a good fit for these wrongly imputed labels does not necessarily lead to an improved $\boldsymbol{\beta}$. In fact, it might lead to worse choices as shown in the results.

Time Complexity

The quadratic programming formulation of ICLS presented in Section 4 allows one to use the standard and constantly improving tools from convex optimization to find the ICLS estimate. Unfortunately one has to go from a convex problem with $d + 1$ variables in the supervised case to a constrained convex problem with U variables for ICLS. For very large U, this may not currently be computationally feasible. Improvements in quadratic programming solvers may help alleviate this problem.

Compared to ICLS, self-learning seems more favorable in terms of computational cost. Self-learning usually converges in a few iterations, where each iteration has the cost of one supervised least squares estimation. In our implementations, however, self-learning and ICLS had similar training times (Figure 6). USM with its simple closed form solution has much lower training times and performs surprisingly well.

Squared Loss

Generally, models used in practice do not directly minimize misclassification error. For computational reasons, often convex surrogate losses, such as the one employed here are minimized. It is therefore interesting to look at the performance of a classifier in terms of these surrogate losses [42]. We have chosen to restrict ourselves to a particular convex loss and attempted to ensure improvement in terms of this chosen loss function.

When we compare the average squared loss on the test set, ICLS, USM and self-learning often seem to offer similar performance. This is quite unlike the results in, for instance [29, 30], where the self-learner often performed much worse in terms of the loss than an approach based on constraining the solution using unlabeled data. While [29, 30] consider a generative classifier, we consider a discriminative classifier, in which case self-learning may be less susceptible to increases in the loss. Self-learning does, however, still increase the loss on some datasets, unlike ICLS.

The peaking phenomenon described in [39, 40] is known to occur for squared loss minimization when we increase the number of labeled samples. Here we find

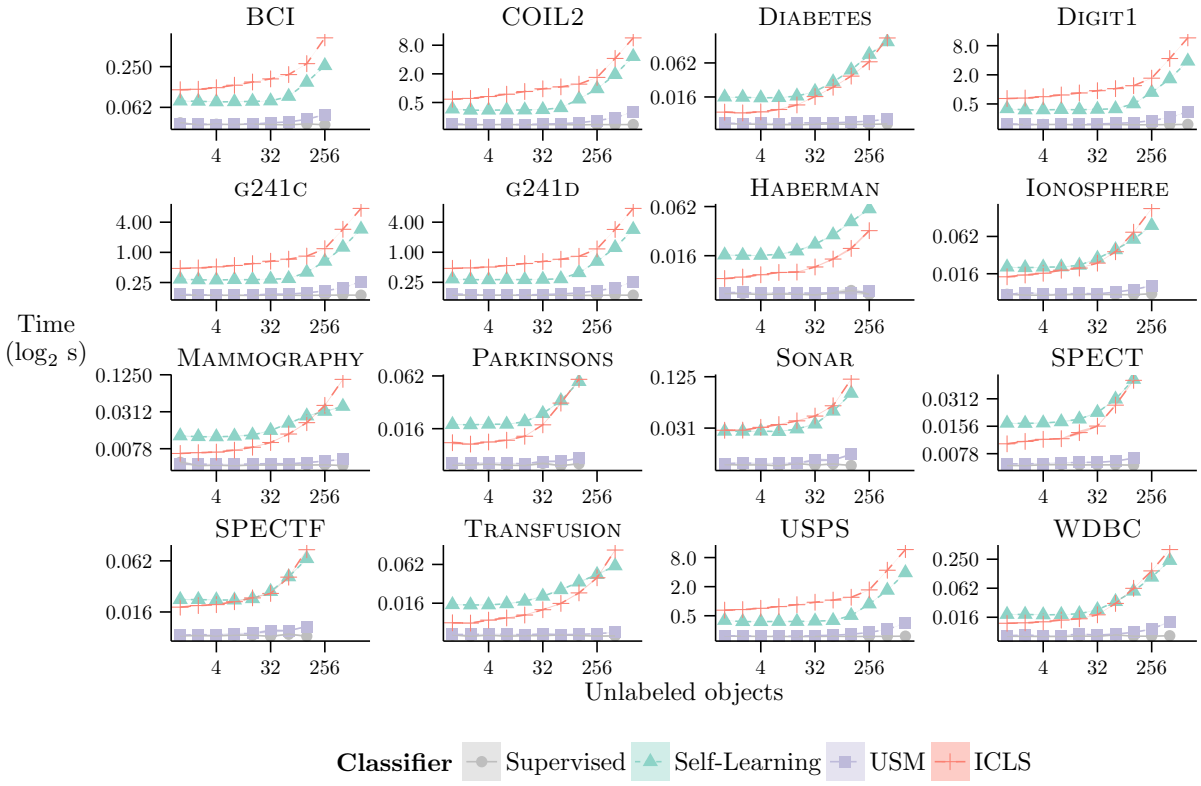

Figure 6: Average Training Time for 1000 repeats. The shaded areas indicate +/− the standard error of the mean.

it also occurs when we change the number of unlabeled samples. It seems that ICLS and USM are more sensitive to this problem than self-learning. As yet, we do not have any explanation for this behavior. Further improvements to the current approach may start by trying to understand this occurrence of peaking.

Other Losses

While the results presented in this work are promising for squared loss, an open question is what other classifiers could benefit from the implicitly constrained approach considered here. Using negative log likelihood as a loss function, for instance, also leads to an interesting implicitly constrained semi-supervised classifier, for instance, in linear discriminant analysis [43].

In the derivation of ICLS, we made use of the closed-form solution given an imputed labeling to derive a quadratic programming problem in terms of the labels. For many loss functions, closed-form solutions do not exist, which prohibits

a straightforward formulation of their implicitly constrained semi-supervised counterparts. Without a supervised closed-form solution one cannot straightaway apply techniques like gradient descent to the parameters as this typically leads to solutions that are outside of the set $\mathcal{C}_\beta$, even if the loss considered is differentiable.

More Constraints

In Figure 1, we illustrate that projecting onto the subset $\mathcal{C}_\beta$ causes improvement as long as a better solution $\hat{\beta}_{oracle}$ than the supervised solution is within $\mathcal{C}_\beta$. A smaller $\mathcal{C}_\beta$ will give a larger improvement, since the projection is going to be closer to $\hat{\beta}_{oracle}$. In the extreme case where only $\hat{\beta}_{oracle}$ forms the subset, this clearly gives a large improvement over supervised learning. It therefore makes sense to think about reducing the size of $\mathcal{C}_\beta$. In the approach presented in this work, however, to ensure a better solution $\hat{\beta}_{oracle}$ than the supervised solution is always within the constraint set with probability $P(\hat{\beta}_{oracle} \in \mathcal{C}_\beta) = 1$, our choice of $\mathcal{C}_\beta$ is conservatively large. It contains elements corresponding to all labelings of the unlabeled points, even extremely unlikely ones.

By excluding unlikely labelings from the subset, the size of $\mathcal{C}_\beta$ may shrink, while the probability that it includes $\hat{\beta}_{oracle}$ remains high. For instance, one might exclude labelings with class priors that are very unlikely to occur, given the class priors that are observed in the labeled data, a strategy which is also employed in Transductive SVMs where it is necessary for it to converge to meaningful local optima. Changes to $\mathcal{C}_\beta$ may, therefore, allow for larger improvements in terms of the risk or classification error, while introducing a small chance of deterioration in performance.

8. Conclusion

This work introduced a new semi-supervised approach to least squares classification. By implicitly considering all possible labelings of the unlabeled objects and choosing the one that minimizes the loss on the labeled observations, we derived a robust classifier with a simple quadratic programming formulation. For this procedure, in the univariate setting with a linear model without intercept, we can prove it never degrades performance in terms of squared loss (Theorem 1). Experimental results indicate that in expectation this robustness also holds in terms of classification error on real datasets. Hence, semi-supervised learning for least squares classification without additional assumptions can lead to improvements over supervised least squares classification both in theory and in practice.

Acknowledgement

Part of this work was funded by project P23 of the Dutch public-private research community COMMIT.

References

References

[1] F. Cozman, I. Cohen, Risks of Semi-Supervised Learning, in: O. Chapelle, B. Schölkopf, A. Zien (Eds.), Semi-Supervised Learning, MIT press, 2006, Ch. 4, pp. 56–72.

[2] F. G. Cozman, I. Cohen, M. C. Cirelo, Semi-Supervised Learning of Mixture Models, in: Proceedings of the Twentieth International Conference on Machine Learning, 2003.

[3] M. Seeger, Learning with labeled and unlabeled data, Tech. rep. (2001).

[4] A. Singh, R. D. Nowak, X. Zhu, Unlabeled data: Now it helps, now it doesn't, in: Advances in Neural Information Processing Systems, 2008, pp. 1513–1520.

[5] T. Hastie, R. Tibshirani, J. H. Friedman, The Elements of Statistical Learning, Spinger, 2001.

[6] T. Poggio, S. Smale, The Mathematics of Learning: Dealing with Data, Notices of the AMS (2003) 537–544.

[7] R. Rifkin, G. Yeo, T. Poggio, Regularized least-squares classification, Nato Science Series Sub Series III Computer and Systems Sciences 190.

[8] J. A. K. Suykens, J. Vandewalle, Least Squares Support Vector Machine Classifiers, Neural Processing Letters 9 (1999) 293–300.

[9] R. Tibshirani, Regression shrinkage and selection via the lasso, Journal of the Royal Statistical Society. Series B 58 (1) (1996) 267–288.

[10] L. Bottou, Large-scale machine learning with stochastic gradient descent, in: Proceedings of COMPSTAT'2010, Springer, 2010, pp. 177–186.

[11] J. H. Krijthe, M. Loog, Implicitly Constrained Semi-Supervised Least Squares Classification, in: E. Fromont, T. D. Bie, M. van Leeuwen (Eds.), 14th International Symposium on Advances in Intelligent Data Analysis XIV, Saint Étienne. France, 2015, pp. 158–169.

[12] O. Chapelle, B. Schölkopf, A. Zien, Semi-supervised learning, MIT press, 2006.

[13] X. Zhu, A. B. Goldberg, Introduction to Semi-Supervised Learning, Vol. 3, Morgan & Claypool, 2009. doi:10.2200/S00196ED1V01Y200906AIM006.

[14] K. Nigam, A. K. McCallum, S. Thrun, T. Mitchell, Text classification from labeled and unlabeled documents using EM, Machine learning 34 (2000) 1–34.

[15] L. Käll, J. D. Canterbury, J. Weston, W. S. Noble, M. J. MacCoss, Semi-supervised learning for peptide identification from shotgun proteomics datasets., Nature methods 4 (11) (2007) 923–925. doi:10.1038/nmeth1113.

[16] M. Shi, B. Zhang, Semi-supervised learning improves gene expression-based prediction of cancer recurrence, Bioinformatics 27 (21) (2011) 3017–3023.

[17] D. Elworthy, Does Baum-Welch re-estimation help taggers?, in: Proceedings of the fourth conference on Applied natural language processing, 1994, pp. 53–58. arXiv:9410012v2.

[18] A. B. Goldberg, X. Zhu, Keepin'it real: semi-supervised learning with realistic tuning, NAACL HLT 2009 Workshop on Semi-supervised Learning for Natural Language Processing.

[19] J. Wang, X. Shen, W. Pan, On Transductive Support Vector Machines, Contemporary Mathematics 443 (2007) 7–19.

[20] G. J. McLachlan, Iterative Reclassification Procedure for Constructing an Asymptotically Optimal Rule of Allocation in Discriminant Analysis, Journal of the American Statistical Association 70 (350) (1975) 365–369.

[21] S. Abney, Understanding the yarowsky algorithm, Computational Linguistics 30 (3) (2004) 365–395.

[22] D. Yarowsky, Unsupervised word sense disambiguation rivaling supervised methods, Proceedings of the 33rd annual meeting on Association for Computational Linguistics (1995) 189–196doi:10.3115/981658.981684.

[23] Y. Grandvalet, Y. Bengio, Semi-supervised learning by entropy minimization, in: L. K. Saul, Y. Weiss, L. Bottou (Eds.), Advances in Neural Information Processing Systems 17, MIT Press, Cambridge, MA, 2005, pp. 529–536.

[24] T. Joachims, Transductive inference for text classification using support vector machines, in: Proceedings of the 16th International Conference on Machine Learning, Morgan Kaufmann Publishers, 1999, pp. 200–209.

[25] K. P. Bennett, A. Demiriz, Semi-supervised support vector machines, in: Advances in Neural Information Processing Systems 11, 1998, pp. 368–374.

[26] V. Sindhwani, S. S. Keerthi, Large scale semi-supervised linear SVMs, in: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval, ACM Press, New York, New York, USA, 2006, p. 477.

[27] R. Collobert, F. Sinz, J. Weston, L. Bottou, Large scale transductive SVMs, Journal of Machine Learning Research 7 (2006) 1687–1712.

[28] J. Wang, X. Shen, Large margin Semi-supervised Learning, Journal of Machine Learning Research 8 (2007) 1867–1891.

[29] M. Loog, Constrained Parameter Estimation for Semi-Supervised Learning: The Case of the Nearest Mean Classifier, in: Proceedings of the 2010 European Conference on Machine learning and Knowledge Discovery in Databases, 2010, pp. 291–304.

[30] M. Loog, A. C. Jensen, Semi-Supervised Nearest Mean Classification through a constrained Log-Likelihood, IEEE Transactions on Neural Networks and Learning Systems 26 (5) (2014) 995 – 1006.

[31] Y.-F. Li, Z.-h. Zhou, Towards making unlabeled data never hurt, in: Proceedings of the 28th International Conference on Machine Learning, 2011, pp. 1081—-1088.

[32] R. J. A. Little, D. B. Rubin, Statistical Analysis with Missing Data, 2002.

[33] M. Healy, M. Westmacott, Missing Values in Experiments Analysed on Automatic Computers, Journal of the Royal Statistical Society 5 (3) (1956) 203–206.

[34] J. P. Shaffer, The Gauss-Markov Theorem and Random Regressors, The American Statistician 45 (4) (1991) 269–273.

[35] B. Fan, Z. Lei, S. Z. Li, Normalized LDA for Semi-supervised Learning, in: 8th IEEE International Conference on Automatic Face & Gesture Recognition, 2008, pp. 1–6. doi:10.1109/AFGR.2008.4813329.

[36] R. H. Byrd, P. Lu, J. Nocedal, C. Zhu, A limited memory algorithm for bound constrained optimization, SIAM Journal on Scientific Computing 16 (5) (1995) 1190–1208.

[37] N. Sokolovska, O. Cappé, F. Yvon, The asymptotics of semi-supervised learning in discriminative probabilistic models, in: W. W. Cohen, A. McCallum, S. T. Roweis (Eds.), Proceedings of the 25th International Conference on Machine Learning, ACM Press, Helsinki, Finland, 2008, pp. 984–991.

[38] K. Bache, M. Lichman, {UCI} Machine Learning Repository (2013).
URL http://archive.ics.uci.edu/ml

[39] M. Opper, W. Kinzel, Statistical Mechanics of Generalization, in: E. Domany, J. L. Hemmen, K. Schulten (Eds.), Models of Neural Networks III, Springer, New York, 1996, pp. 151–209. doi:10.1007/978-1-4612-0723-8{_}5.

[40] S. Raudys, R. P. Duin, Expected classification error of the Fisher linear classifier with pseudo-inverse covariance matrix, Pattern Recognition Letters 19 (5-6) (1998) 385–392. doi:10.1016/S0167-8655(98)00016-6.

[41] M. Skurichina, R. P. W. Duin, Regularisation of Linear Classifiers by Adding Redundant Features, Pattern Analysis & Applications 2 (1) (1999) 44–52. doi:10.1007/s100440050013.

[42] M. Loog, J. H. Krijthe, A. C. Jensen, On Measuring and Quantifying Performance: Error Rates, Surrogate Loss, and an Example in SSL, in: C. H. Chen (Ed.), Handbook of Pattern Recognition and Computer Vision, 5th Edition, World Scientific, 2016, Ch. 1.3.

[43] J. H. Krijthe, M. Loog, Implicitly Constrained Semi-Supervised Linear Discriminant Analysis, in: International Conference on Pattern Recognition, Stockholm, 2014, pp. 3762–3767.